

Improved Particle Filters for Vehicle Localisation

Kira Kempinska

*Department of Security and Crime Science
University College London
London, WC1E 6BT*

KIRA.KOWALSKA.13@UCL.AC.UK

John Shawe-Taylor

*Department of Computer Science
University College London
London, WC1E 6BT*

J.SHAWE-TAYLOR@UCL.AC.UK

Abstract

The ability to track a moving vehicle is of crucial importance in numerous applications. The task has often been approached by the importance sampling technique of particle filters due to its ability to model non-linear and non-Gaussian dynamics, of which a vehicle travelling on a road network is a good example. Particle filters perform poorly when observations are highly informative. In this paper, we address this problem by proposing particle filters that sample around the most recent observation. The proposal leads to an order of magnitude improvement in accuracy and efficiency over conventional particle filters, especially when observations are infrequent but low-noise.

1. Introduction

Tracking a moving vehicle is a central and difficult problem arising in different contexts ranging from military applications to robotics (Thrun, 2002; Gordon et al., 2002). It consists of computing the best estimate of the vehicle's trajectory based on noisy sensor measurements. In this paper, we are interested in vehicle tracking when the road network is known.

Several strategies have been developed to track a vehicle on a road network (Lou et al., 2009; Chawathe, 2007; Wenk et al., 2006; Alt et al., 2003; Huabei Yin and Wolfson; Pink and Hummel, 2008). We focus on the particle filter method (Gordon et al., 1993). The method has had numerous successes in this area due to its flexibility to handle cases where the dynamic and observation models are non-linear and/or non-Gaussian. It is an importance sampling technique that approximates the target distribution by sampling from a series of intermediate proposal distributions.

Critically, in common with any important sampling method, the performance of particle filters is strongly dependent on the choice of the proposal distribution. If the proposal is not well matched to the target distribution, then the method produces samples that have low effective sample size and, as a result, it requires a prohibitively large number of particles to represent the target distribution accurately. The problem typically arises under highly informative observation regimes, in which the current observation provides significant information about the current state but the state dynamics are weak.

The particle filter community has developed various approaches to mitigate the deficiency. One approach attaches a post-sampling step that moves particles sampled from the proposal distribution towards the target distribution using Markov Chain Monte Carlo moves (Andrieu et al., 2010; Fox et al., 2001; Van Der Merwe et al., 2000) or by solving partial differential equations (Li and Coates, 2016; Daum et al., 2010; Daum and Huang, 2007; Khan and Ulmke, 2014). An alternative approach improves the proposal distribution by giving it additional information about the current (Montemerlo and Thrun, 2007; Kong et al., 1994; Liu and Chen, 1998) or even future observations (Lin et al., 2013) or their approximations (Pitt and Shephard, 1999). Several authors considered conditioning the proposal distribution on the current observation only (Lin et al., 2005; Fox et al., 2001). The

approaches successfully increased the effective sample size, but, often at the cost of high computational complexity or analytical intractability. The construction of good, but also computationally efficient proposal distributions is still an open research question.

In this paper, we propose an improved particle sampling scheme that is both computationally efficient and mathematically robust. The approach generates proposals based on the current sensor observation only, leading to good alignment between the proposal and the target distribution even with a small sample size. It converges to the desired target distribution at faster rates than standard particle filters, especially when observations are highly informative, e.g. infrequent but low-noise. It is easy to implement and avoids the computational and analytical complexity of the discussed alternatives with other proposal distributions or post-sampling moves. It also presents a simpler approach to sample weighing than those previously proposed with the same proposal distribution (Lin et al., 2005; Fox et al., 2001).

The paper is structured as follows. We present the problem statement in Section 2, followed by a description of the standard particle filters in Section 3. We introduce the proposed particle filters method in Section 4. We outline the application of the method to vehicle tracking in Section 5 and present results in Section 6. We conclude by summarising the paper’s contributions in Section 7.

2. Problem statement

The key idea of particle filters is to estimate the marginal posterior distribution $p(x_t | z_{0:t})$ where x_t is the state of the system at time t and $z_{0:t} = \{z_0, \dots, z_t\}$ is a sequence of measurements collected up to time step t . We call the posterior the *belief* and use the following notation

$$Bel(x_t) = p(x_t | z_{0:t}) \quad (1)$$

In the context of vehicle tracking, the belief is our estimate of the vehicle position at time t given all measurements collected until then. The measurements include *GPS readings* and *controls*, which carry information about vehicle motion between consecutive timestamps. Denoting a GPS reading at time t by y_t and a control in the time interval $(t-1; t]$ by u_{t-1} , we have

$$Bel(x_t) = p(x_t | y_{0:t}, u_{0:t-1}) \quad (2)$$

Particle filters estimate $Bel(x_t)$ recursively. In order to arrive at a recursive equation, we note we can use Bayes rule to decompose Equation 2 to

$$Bel(x_t) = \frac{p(y_t | x_t, y_{0:t-1}, u_{0:t-1})p(x_t | y_{0:t-1}, u_{0:t-1})}{p(y_t | y_{0:t-1}, u_{0:t-1})} \quad (3)$$

The underlying assumption of particle filters is that the system follows the *Markov assumption*, that is, measurements y_t are conditionally independent of past measurements and controls given knowledge of the state x_t :

$$p(y_t | x_t, y_{0:t-1}, u_{0:t-1}) = p(y_t | x_t)$$

This conveniently simplifies Equation 3 to

$$Bel(x_t) = \frac{p(y_t | x_t)p(x_t | y_{0:t-1}, u_{0:t-1})}{p(y_t | y_{0:t-1}, u_{0:t-1})}$$

We integrate out the position at x_{t-1} in order to arrive at the following recursive form

$$Bel(x_t) = \frac{p(y_t | x_t)}{p(y_t | y_{0:t-1}, u_{0:t-1})} \int p(x_t | x_{t-1}, y_{0:t-1}, u_{0:t-1})p(x_{t-1} | y_{0:t-1}, u_{0:t-1})dx_{t-1}$$

which can be simplified again using *Markov assumption* by noting that:

$$p(x_t \mid x_{t-1}, y_{0:t-1}, u_{0:t-1}) = p(x_t \mid x_{t-1}, u_{t-1})$$

Finally, we arrive at a recursive estimator known as *Bayes filter*:

$$\begin{aligned} Bel(x_t) &= \frac{p(y_t \mid x_t)}{p(y_t \mid y_{0:t-1}, u_{0:t-1})} \int p(x_t \mid x_{t-1}, u_{t-1}) Bel(x_{t-1}) dx_{t-1} \\ &= \eta p(y_t \mid x_t) \int p(x_t \mid x_{t-1}, u_{t-1}) Bel(x_{t-1}) dx_{t-1} \end{aligned} \quad (4)$$

where η is a normalising constant. The *Bayes filter* equation is the basis for particle filters and the improved particle filters that we propose in this paper.

3. Particle Filters

Particle filters approximate the belief $Bel(x)$ by a set of m weighted samples distributed according to $Bel(x)$:

$$Bel(x) = \{x^{(i)}, w^{(i)}\}_{i=1, \dots, m}$$

where each $x^{(i)}$ is a sample (a state) and $w^{(i)}$ are non-negative weights called *importance factors* that determine the importance of each sample.

The particle filters method operates recursively. It begins by generating m samples $x_0^{(i)}$ from the initialisation distribution $Bel(x_0) = p(x_0)$ and annotates them by the uniform importance factor $1/m$. Subsequently, it estimates $Bel(x_t)$ at any future timestamp t by performing a three-step recursive update, computing the expression in Equation 4 *from the right to the left*.

for $k = 1, \dots, m$:

1. Sample a state x_{t-1} by drawing a random $x_{t-1}^{(i)}$ from the sample set representing $Bel(x_{t-1})$ according to the distribution defined through the importance factors $w_{t-1}^{(i)}$.
2. Use the sample $x_{t-1}^{(i)}$ and the control u_{t-1} to generate a sample $x_t^{(j)}$ according to the so-called *transition probability* $p(x_t \mid x_{t-1}, u_{t-1})$.
3. Finally, use the observation y_t to weigh the sample $x_t^{(j)}$ by the non-normalized importance factor given by the so-called *observation probability* $p(y_t \mid x_t^{(j)})$, the likelihood of the sample $x_t^{(j)}$ given the observation y_t .

Further below, it will be important to notice that the particle filters method is, in fact, an importance sampling scheme. It approximates $Bel(x_t)$ using a proposal distribution given by

$$Q = p(x_t \mid x_{t-1}, u_{t-1}) Bel(x_{t-1}) \quad (5)$$

The proposal approximates the desired posterior

$$P = \frac{p(y_t \mid x_t) p(x_t \mid x_{t-1}, u_{t-1}) Bel(x_{t-1})}{p(y_t \mid y_{0:t-1}, u_{0:t-1})} \quad (6)$$

Consequently, the importance factors are given by the quotient

$$\begin{aligned} \frac{P}{Q} &= [p(x_t | x_{t-1}, u_{t-1}) Bel(x_{t-1})]^{-1} \frac{p(y_t | x_t) p(x_t | x_{t-1}, u_{t-1}) Bel(x_{t-1})}{p(y_t | y_{0:t-1}, u_{0:t-1})} \\ &\propto p(y_t | x_t) \end{aligned} \quad (7)$$

4. Improved Particle Filters

We propose an improved particle sampling scheme in which x_t are sampled directly around the most recent observation y_t according to the proposal distribution:

$$Q_{new} = \frac{p(y_t | x_t)}{\pi(y_t)} \quad \text{with} \quad \pi(y_t) = \int p(y_t | x_t) dx_t \quad (8)$$

This new proposal distribution possesses orthogonal strengths to the one in Equation 5, in that it generates samples that are highly consistent with the most recent sensor measurement but ignorant of past measurements and controls. As such, we expect it to outperform conventional particle filters in systems where the current observation provides more information about the current state than the underlying state dynamics.

The importance factors for these samples are again calculated by the quotient:

$$\begin{aligned} \frac{P}{Q_{new}} &= \left[\frac{p(y_t | x_t)}{\pi(y_t)} \right]^{-1} \frac{p(y_t | x_t) p(x_t | x_{t-1}, u_{t-1}) Bel(x_{t-1})}{p(y_t | y_{0:t-1}, u_{0:t-1})} \\ &= \frac{p(x_t | x_{t-1}, u_{t-1}) Bel(x_{t-1}) \pi(y_t)}{p(y_t | y_{0:t-1}, u_{0:t-1})} \\ &\propto p(x_t | x_{t-1}, u_{t-1}) Bel(x_{t-1}) \end{aligned} \quad (9)$$

Since $Bel(x_{t-1})$ is represented by a set of samples $x_{t-1}^{(i)}$ weighted by importance factors $w_{t-1}^{(i)}$, the (non-normalised) importance factor for any sample $x_t^{(j)}$ can be approximated by

$$\sum_{i=1}^m p(x_t^{(j)} | x_{t-1}^{(i)}, u_{t-1}) w_{t-1}^{(i)} \quad (10)$$

The importance factor reflects the likelihood of the sample given *past* measurements and controls. This is orthogonal to the previous definition in Equation 7, where it depends on the *current* measurement only.

Overall, the proposed sampling scheme changes how data and controls are used in belief estimation: the current measurement is now used for sampling (instead of weighing); past measurements and controls are used for calculating importance factors (instead of sampling).

The scheme is implemented recursively. It initialises $Bel(x_0)$ by generating m samples around the first observation y_0 according to the observation probability $p(y_t | x_t)$. The samples are assigned the uniform importance factor of $1/m$. Subsequently, it estimates $Bel(x_t)$ at timestamps $t > 0$ using a two-step recursive update:

for $k = 1, \dots, m$:

1. Generate a sample $x_t^{(i)}$ according to the observation probability $p(y_t | x_t)$.
2. Use the sample set representing $Bel(x_{t-1})$ to weight the sample $x_t^{(i)}$ by the importance factor in Equation 10 , the likelihood of the sample $x_t^{(i)}$ given past measurements and controls.

In the context of vehicle tracking, the estimates of $Bel(x_t)$ approximate the vehicle *position* at time t . If instead of the single-time approximation, you are interested in finding the most likely *trajectory* that the vehicle traversed until time t , it can be computed via the following dynamic programming routine. It corresponds to finding the sequence $x_{0:t}$ that maximises the posterior $p(x_{0:t} | y_{0:t}, u_{0:t-1})$.

1. Choose a sample $x_t^{(i)}$ from the sample set representing $Bel(x_t)$ that has the highest importance factor $w_t^{(i)}$.
2. Use the sample $x_t^{(i)}$ to find a preceding sample $x_{t-1}^{(j)}$ from $Bel(x_{t-1})$ that maximises $p(x_t^{(i)} | x_{t-1}^{(j)}, u_{t-1})w_{t-1}^{(j)}$, i.e. is the most likely preceding state. Repeat this step until you reach $t = 0$.

5. Application to Vehicle Tracking

5.1 Data

We tested the improved particle filters method on a GPS trajectory of a police patrol vehicle during its night shift (9am to 7am) in the London Borough of Camden on February 9th 2015. The dataset contains 4,800 GPS points that were emitted roughly every second when moving. It was acquired for research purposes as part of the "Crime, Policing and Citizenship" project¹.

5.2 Implementation

In order to apply the improved particle filters to vehicle tracking, we need to specify the form of the observation probability $p(y_t | x_t)$ and the transition probability $p(x_t | x_{t-1}, u_{t-1})$. Their forms depend on the vehicle's dynamics and the type of sensor used for localisation (a GPS receiver in our case). The distributions are time-invariant; hence we will omit the time index t in the following derivations.

5.2.1 OBSERVATION PROBABILITY

Definition We model the conditional probability $p(y | x)$ of observing a GPS point \mathbf{y} , represented by its easting and northing coordinates:

$$\mathbf{y} = \begin{pmatrix} y_e \\ y_n \end{pmatrix} \quad (11)$$

as a two-dimensional Gaussian distribution

$$\mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$$

with the mean vector $\boldsymbol{\mu}$ representing the true vehicle position \mathbf{x}

$$\boldsymbol{\mu} = \mathbf{x} = \begin{pmatrix} x_e \\ x_n \end{pmatrix} \quad (12)$$

1. UCL Crime Policing and Citizenship: <http://www.ucl.ac.uk/cpc/>

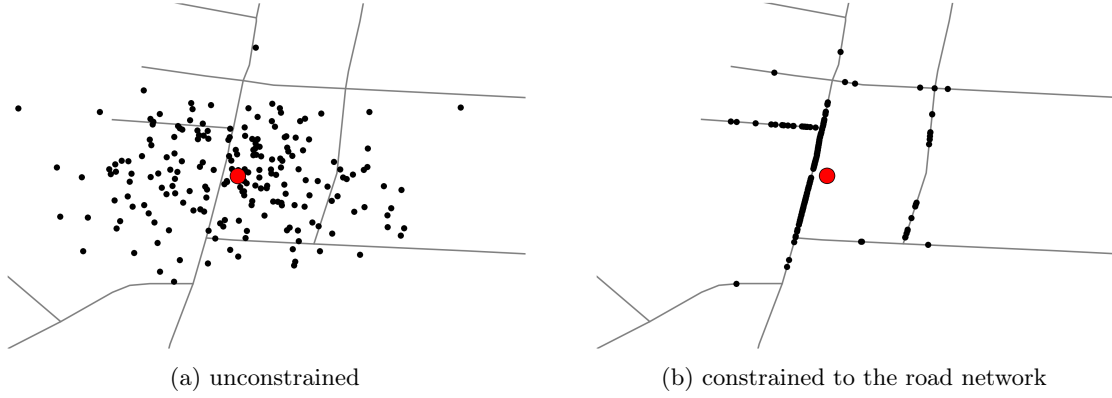


Figure 1: Sampling proposal positions around a GPS point (red).

and the covariance Σ that is constant across space, i.e. *isotropic* covariance

$$\Sigma = \begin{bmatrix} \Sigma_{ee} & \Sigma_{en} \\ \Sigma_{ne} & \Sigma_{nn} \end{bmatrix} = \begin{bmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{bmatrix} \quad (13)$$

This representation of $p(y | x)$ reflects our expectation that GPS observations are normally distributed around the true vehicle positions.

Proposal Generation In the proposed method, we use the observation probability $p(y | x)$ to sample possible vehicle positions on the road network (see Equation 8). In order to efficiently generate samples on the road network (as shown in Figure 1b), we want to analytically project the two-dimensional $p(y | x)$ onto individual road segments.

We begin with the general form of a two-dimensional Gaussian distribution for $p(y | x)$

$$p(y | x) = \mathcal{N}(\mathbf{y} | \boldsymbol{\mu}, \Sigma) = \frac{1}{2\pi|\Sigma|^{1/2}} \exp \left\{ -\frac{1}{2}(\mathbf{y} - \boldsymbol{\mu})^T \Sigma^{-1}(\mathbf{y} - \boldsymbol{\mu}) \right\} \quad (14)$$

We precompute the inverse of the covariance matrix

$$\Sigma^{-1} = \frac{1}{\sigma^4} \begin{bmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{bmatrix} = \begin{bmatrix} \sigma^{-2} & 0 \\ 0 & \sigma^{-2} \end{bmatrix}$$

and use it together with the partitioning (11), (12), and (13) to rewrite (14) as

$$\begin{aligned} p(y | x) &= \frac{1}{2\pi\sigma^2} \exp \left\{ -\frac{1}{2} \left[\frac{(y_e - \mu_e)^2}{\sigma^2} + \frac{(y_n - \mu_n)^2}{\sigma^2} \right] \right\} \\ &= \frac{1}{(2\pi\sigma^2)^{1/2}} \exp \left\{ -\frac{1}{2\sigma^2} (y_e - \mu_e)^2 \right\} \times \frac{1}{(2\pi\sigma^2)^{1/2}} \exp \left\{ -\frac{1}{2\sigma^2} (y_n - \mu_n)^2 \right\} \\ &= \mathcal{N}(y_e | \mu_e, \sigma) \times \mathcal{N}(y_n | \mu_n, \sigma) \end{aligned} \quad (15)$$

We successfully factor $p(y | x)$ into a product of two Gaussian distributions along the *easting* and *northing* directions due to the isotropic properties of the covariance matrix in (13). In fact, the factorisation of $p(y | x)$ holds for any other orthogonal coordinate system. Therefore, we replace the easting-northing coordinates with orthogonal distances from x dictated by the road segment that x is on: a (distance *to* the road segment), b (distance *along* the road segment).

Under the new coordinate system x and y are partitioned as

$$x = \begin{pmatrix} x_a \\ x_b \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \quad y = \begin{pmatrix} y_a \\ y_b \end{pmatrix}$$

and $p(y | x)$ becomes

$$\begin{aligned} p(y | x) &= \mathcal{N}(y_a | \mu_a, \sigma) \times \mathcal{N}(y_b | \mu_b, \sigma) \\ &= \mathcal{N}(y_a | 0, \sigma) \times \mathcal{N}(y_b | 0, \sigma) \end{aligned} \tag{16}$$

The above definition enables us to generate proposals x in accordance with the observation model $p(y | x)$ (as specified in Equation 8):

1. Firstly, sampling a road segment that x is on such that $y_a \sim \mathcal{N}(0, \sigma)$
2. Secondly, sampling the position of x along the segment such that $y_b \sim \mathcal{N}(0, \sigma)$

5.2.2 TRANSITION PROBABILITY

We set the transition probability $p(x_t | x_{t-1}, u_{t-1})$ to be a linear estimate equal to the Cartesian distance between GPS points x_{t-1} and x_t (the control u_t) plus an additive Gaussian noise. This is a simplistic assumption that could be further explored, however, it is not the focus of this paper.

5.3 Validation

In the absence of the ground truth about vehicle positions at any point in time, we propose a validation framework based on the well-established technique of cross-validation (Barber, 2012). We remove every 10th GPS points from the available GPS trajectory. We then infer the path taken by the vehicle given the incomplete trajectory and the road network. Finally, we measure the distance between each removed point and the inferred path. The distances across all removed points form the distribution of the prediction error.

6. Results

A series of tests was conducted to elucidate the difference between the standard and the proposed particle filters. We found that the modified proposal distribution consistently outperforms conventional particle filters in terms of accuracy. As expected, largest gains in accuracy are observed on datasets with long sampling intervals as their observations are infrequent and hence become highly informative. Figure 2 plots the prediction error (in meters) of both algorithms for different sampling intervals and levels of sensor noise, using $m = 10$ samples only. It shows that the proposed method has lower *median* error across all examined sampling rates and sensor noise levels, as well as much lower error *variation*.

We evaluated the ability of both methods to track a vehicle over time. When they fail to track a vehicle, it means that all positions that they propose are completely unlikely given sensor data, i.e. particle weights sum up to zero. The standard particle filter basically fails when sensor measurements are infrequent (with $m = 10$ samples). Figure 3a shows that it is unable to track the vehicle nearly 70% of the time when the sampling interval increases to one minute. In the same scenario, the proposed method gives excellent results that show little variation to changes to sampling intervals.

On the contrary, the proposed method fails to track when sensors are very noisy. Although it shows high accuracy (see Figure 2b), it is prone to high failure rates as the level of sensor noise increases (Figure 3b). This weakness reflects the orthogonal limitations of the two approaches: our method generates samples that are highly consistent with the most recent measurement (which makes it sensitive to sensor noise), whereas the conventional approach samples in accordance with past measurements (inefficient when sampling rates are low).




Figure 2: Accuracy of the improved particle filters (red) and the standard particle filters (blue) on GPS data with varied sampling rate and sensor noise, represented as 25th, 50th and 75th percentiles of prediction errors.




Figure 3: Percentage of time the improved particle filters (red) and the conventional particle filters (blue) lost track of the position of the vehicle as a function of the GPS sampling rate and the sensor noise.




Figure 4: Accuracy and robustness of the improved particle filters (red) and the conventional particle filters (blue) as a function of the number of samples used. Accuracy is shown as the 25th, 50th and 75th percentiles of prediction errors.




Figure 5: Accuracy and robustness of the standard particle filters as the number of samples is increased to very large values. Accuracy is shown as the 25th, 50th and 75th percentiles of prediction errors.

Finally, we tested the sensitivity of the proposed method to the number of samples used. Figure 4 shows comparative results on GPS data with the sampling interval of 70 seconds. The proposed method yields significantly better results, both in terms of accuracy and robustness to failure. When only $m = 10$ samples are used, it reduces the estimation error by almost 10 meters and the percentage of failure by as much as 68%. The performance is further improved when more samples are used, but the gain is small compared to the conventional particle filters. In fact, the proposed method with $m = 100$ samples is more accurate and robust than the conventional particle filters with as many as $m = 10000$ samples (see Figure 5). Therefore, it can be reliably used with a small number of samples, making it highly computationally efficient.

7. Conclusions

This paper describes a modified particle filters method that shows uniformly superior accuracy to the conventional particle filters. The improved algorithm utilizes a different proposal distribution which uses only the most recent observation in the position prediction process. In doing so, it makes more efficient use of the particles, particularly in situations in which the transition noise is high in relation to the observation noise.

The main contribution of the paper is the proposal distribution itself and the derivation of the associated importance weights that guarantees convergence to the same posterior distribution as the standard particle filters. An important contribution is also the projection of a two-dimensional Gaussian onto a network of roads, which enables efficient sampling on the road network from a spatial Gaussian.

The theoretical contributions are complemented by experimental results of vehicle tracking using a police GPS dataset. The new algorithm is consistently more accurate than the standard particle filters, with largest gains in accuracy on sparse GPS data. It requires much fewer samples to yield good performance. In fact, as few as fifty samples are sufficient to outperform the standard method with 10,000 particles in terms of accuracy and proneness to failure. We believe that our results illustrate that particle filters can be radically improved if one carefully chooses a proposal distribution, such that it extracts the most information from the available data.

Acknowledgements

This work is part of the project - Crime, Policing and Citizenship (CPC): Space-Time Interactions of Dynamic Networks (www.ucl.ac.uk/cpc), supported by the UK Engineering and Physical Sciences Research Council (EP/J004197/1). The data provided by Metropolitan Police Service (London) is greatly appreciated.

We would also like to show our gratitude to Dr Simon Julier for very helpful discussions during the course of this research.

References

- Helmut Alt, Alon Efrat, Günter Rote, and Carola Wenk. Matching planar maps. *Journal of Algorithms*, 49(2):262–283, 2003. ISSN 01966774. doi: 10.1016/S0196-6774(03)00085-3.
- Christophe Andrieu, Arnaud Doucet, and Roman Holenstein. Particle Markov chain Monte Carlo methods. *J. R. Statist. Soc. B*, 72(3):269–342, 2010.
- David Barber. *Bayesian Reasoning and Machine Learning*. Cambridge University Press, 2012. ISBN 0521518148.
- Sudarshan S. Chawathe. Segment-Based Map Matching. In *2007 IEEE Intelligent Vehicles Symposium*, pages 1190–1197. IEEE, jun 2007. ISBN 1-4244-1067-3. doi: 10.1109/IVS.2007.4290280.
- Fred Daum and Jim Huang. Nonlinear filters with log-homotopy. In Oliver E. Drummond and Richard D. Teichgraeber, editors, *Proceedings of SPIE*, sep 2007. doi: 10.1117/12.725684.
- Fred Daum, Jim Huang, and Arjang Noushin. Exact particle flow for nonlinear filters. In *SPIE Defense, Security, and Sensing*, pages 769704–769704. International Society for Optics and Photonics, 2010.
- Dieter Fox, Sebastian Thrun, Wolfram Burgard, and Frank Dellaert. Particle Filters for Mobile Robot Localization. In *Sequential Monte Carlo Methods in Practice*, pages 401–428. Springer New York, New York, NY, 2001. doi: 10.1007/978-1-4757-3437-9_19.

- Neil J. Gordon, Simon Maskell, and Thiagalingam Kirubarajan. Efficient particle filters for joint tracking and classification. In Oliver E. Drummond, editor, *Signal and Data Processing of Small Targets*, Orleando, 2002. International Society for Optics and Photonics. doi: 10.1117/12.478524.
- N.J. Gordon, D.J. Salmond, and A.F.M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F Radar and Signal Processing*, 140(2):107, 1993. ISSN 0956375X. doi: 10.1049/ip-f-2.1993.0015.
- Huabei Yin and O. Wolfson. A weight-based map matching method in moving objects databases. In *Proceedings. 16th International Conference on Scientific and Statistical Database Management, 2004.*, pages 437–438. IEEE. ISBN 0-7695-2146-0. doi: 10.1109/SSDM.2004.1311248.
- Muhammad Altamash Khan and Martin Ulmke. Non-linear and non-Gaussian state estimation using log-homotopy based particle flow filters. In *2014 Sensor Data Fusion: Trends, Solutions, Applications (SDF)*, pages 1–6. IEEE, oct 2014. ISBN 978-1-4799-7387-3. doi: 10.1109/SDF.2014.6954715.
- Augustine Kong, Jun S Liu, and Wing Hung Wong. Sequential Imputations and Bayesian Missing Data Problems. *Journal of the American Statistical Association*, 89(425):278–288, 1994.
- Yunpeng Li and Mark Coates. Particle Filtering with Invertible Particle Flow. jul 2016.
- Ming Lin, Rong Chen, and Jun S. Liu. Lookahead Strategies for Sequential Monte Carlo. *Statistical Science*, 28(1):69–94, feb 2013. ISSN 0883-4237. doi: 10.1214/12-STS401.
- Ming T Lin, Junni L Zhang, Qiansheng Cheng, and Rong Chen. Independent Particle Filters. *Journal of the American Statistical Association*, 100(472):1412–1421, dec 2005. ISSN 0162-1459. doi: 10.1198/016214505000000349.
- Jun S. Liu and Rong Chen. Sequential Monte Carlo Methods for Dynamic Systems. *Journal of the American Statistical Association*, 93(443):1032, sep 1998. ISSN 01621459. doi: 10.2307/2669847.
- Yin Lou, Chengyang Zhang, Yu Zheng, Xing Xie, Wei Wang, and Yan Huang. Map-matching for low-sampling-rate GPS trajectories. In *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems - GIS '09*, page 352, New York, New York, USA, nov 2009. ACM Press. ISBN 9781605586496. doi: 10.1145/1653771.1653820.
- Michael Montemerlo and Sebastian Thrun. FastSLAM 2.0. In *FastSLAM*, pages 63–90. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007. doi: 10.1007/978-3-540-46402-0_4.
- Oliver Pink and Britta Hummel. A statistical approach to map matching using road network geometry, topology and vehicular motion constraints. In *2008 11th International IEEE Conference on Intelligent Transportation Systems*, pages 862–867. IEEE, oct 2008. ISBN 978-1-4244-2111-4. doi: 10.1109/ITSC.2008.4732697.
- Michael K. Pitt and Neil Shephard. Filtering via Simulation: Auxiliary Particle Filters. *Journal of the American Statistical Association*, 94(446):590–599, jun 1999. ISSN 0162-1459. doi: 10.1080/01621459.1999.10474153.
- Sebastian Thrun. Particle filters in robotics. In *Proceedings of Uncertainty in AI*, pages 511–518. Morgan Kaufmann Publishers Inc., aug 2002. ISBN 1-55860-897-4.
- Rudolph Van Der Merwe, N Ando De Freitas, and Eric Wan. The Unscented Particle Filter. *Advances in Neural Information Processing Systems*, pages 584–590, 2000.
- Carola Wenk, Randall Salas, and Dieter Pfoser. Addressing the Need for Map-Matching Speed: Localizing Global Curve-Matching Algorithms. In *Proceedings of the 18th international Conference on Scientific and Statistical Database Management*, 2006.